DOCUMENT RESUME

ED 405 092 PS 025 028

AUTHOR Halford, Graeme S.
TITLE Relational Knowledge in Higher Cognitive Processes.
PUB DATE Aug 96
NOTE 17p.; Paper presented at the Biennial Meeting of the International Society for the Study of Behavioral Development (14th, Quebec City, Quebec, Canada, August 12-16, 1996).
PUB TYPE Speeches/Conference Papers (150) -- Information Analyses (070)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS Associative Learning; *Children; Classification; *Cognitive Processes; Cognitive Structures; Deduction; Foreign Countries; Fractions; Planning; Ratios (Mathematics); *Schemata (Cognition); Thinking Skills
IDENTIFIERS Analogies; Class Inclusion; Neural Net Models; *Relational Concepts; *Relational Learning; Transitivity

ABSTRACT

Explicit representation of relations plays some role in virtually all higher cognitive processes, but relational knowledge has seldom been investigated systematically. This paper considers how relational knowledge is involved in some tasks that have been important to cognitive development, including transitivity, the balance scale, classification hierarchies, and deductive reasoning. It is proposed that important properties of these tasks are due to the fact that they entail processing relations. The nature of relational knowledge, its implementation in neural nets, and its role in higher cognition are examined. The paper outlines the properties of relational knowledge and distinguishes it from more primitive processes such as traditional associationism. Relational knowledge is described as flexible, explicit, and can be organized into complex structures such as lists, trees, and propositional networks. Relational knowledge is central to mechanisms that are basic to human reasoning, such as analogy and planning. The properties of relational knowledge are obtained at the cost of higher processing loads. Empirical criteria for relational knowledge are also indicated. (Author/KDFB)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)
X This document has been reproduced as received from the person or organization originating it.
☐ Minor changes have been made to improve reproduction quality.
• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

ED 405 092

Relational Knowledge in Higher Cognitive Processes

Graeme S. Halford

University of Queensland

This paper considers how relational knowledge is involved in some tasks that have been important to cognitive development, including transitivity, the balance scale, classification hierarchies, and deductive reasoning. It is proposed that important properties of these tasks are due to the fact that they entail processing relations. The nature of relational knowledge, its implementation in neural nets, and its role in higher cognition, will be examined. The paper outlines the properties of relational knowledge, and distinguishes it from more primitive processes (e.g. traditional associationism). Relational knowledge is flexible, explicit, and can be organized into complex structures such as lists, trees and propositional networks. It is central to mechanisms that are basic to human reasoning, such as analogy and planning. However its properties are obtained at the cost of higher processing loads. Empirical criteria for relational knowledge are also indicated.

Paper presented in symposium entitled Relational Knowledge in Higher Cognitive Processes at the XIVth Biennial Meeting of the International Society for the Study of Behavioral Development, Quebec City, August 12-16, 1996.

PERMISSION TO REPRODUCE AND
DISSEMINATE THIS MATERIAL
HAS BEEN GRANTED BY
G.S. Halford

TO THE EDUCATIONAL RESOURCES
INFORMATION CENTER (ERIC)

PS 025028

BEST COPY AVAILABLE

Explicit representation of relations plays some kind of role in virtually all higher cognitive processes, and has been important historically, yet Psychologists have made little attempt to investigate relational knowledge systematically. This is incongruous in itself, and contrasts with disciplines like computer science where relational data bases are a major field of work. This symposium is an attempt to draw attention to this question, and to show how analyses of relations in cognitive tasks can shed light on some familiar problems.

Relations are basic to mathematics, as illustrated by the fact that the familiar arithmetic operations addition and multiplication are ternary relations, and a mathematical function is a special case of a relation. Understanding mathematics depends in part on representing some of the higher-order relations that link mathematical concepts. Development of procedures and strategies is guided by understanding of the relations that underlie tasks.

Some of the tasks that have mysteriously caused difficulty for young children are found to entail complex relations that impose high processing loads. Some of these, such as transitivity and class inclusion, originate in the Piagetian tradition, but others such as concept of mind originated in contemporary developmental research, while still others such as conditional discrimination originated in a behaviouristic learning tradition. I want to propose that many properties of these tasks, including the underlying similarity between them, can be recognised by analysing the relations they entail.

In order to avoid misunderstanding I would like to emphasise that when we analyse underlying relations in tasks, and thereby recognise commonalities between different domains, we also recognise there is knowledge that is specific to each domain. I do not want to deny the importance of domain knowledge, but the underlying relations in tasks have received less attention than they deserve.

Rules

setting condition	object-attribute	person's percept
no-filter	white	white
blue-filter	white	blue

Relational representation

Appearance-reality(<condition/event>,<object-attribute>,<percept>)

Figure 1. Simple appearance-reality task analysed as rules and relations.

We will first consider the concept of mind task, which has been analysed both in my own work and in a paper by Frye, Zelazo and Palfai (1995). One of the simplest tasks entails a white fish which can be viewed with or without a blue filter. The child is asked how the fish appears when viewed through the blue filter, and also how the object is "really and truly". Some of the younger children seem to be able to answer solely in terms of the way the object appears, or in terms of the way it really is, but cannot recognise both appearance and reality. Therefore John Flavell suggested that they cannot handle two ways of representing an object. The rules are shown in Figure 1.

According to the analysis of Frye et al. (1995) the filter would be considered a setting condition, and adds an extra level to a hierarchical representation. In my own analysis (Halford, 1993; 1996) there are two relations between the person's percept and the object, and these are integrated into a higher-order relation. By either analysis the essential idea is that the relation between an object and a person's perception of that object is conditional on a third variable which, in this case, is the presence or absence of the filter. The relational representation is also shown in the Figure. This does not invalidate the other analyses, but the formal expression of relations provides the most general representation of task structure, and permits comparison across all contexts, irrespective of whether they are hierarchical or not.

APPEARANCE-REALITY

Setting condition		object attribute		person's percept
no-filter		white	→	white
blue-filter	:	white	→	blue

CONDITIONAL DISCRIMINATION

background	cue		response
black	triangle	→	R+
black	square	→	R-
white	triangle	→	R-
white	square	→	R+

CONFIGURAL LEARNING

Configuration		response
Configuration triangle/black	→	R+
Configuration square/black	→	R-
Configuration triangle/white	→	R-
Configuration square/white	→	R+

ISOMORPHIC CONDITIONAL DISCRIMINATION

background	cue		response
green	circle	→	R+
blue	cross	→	?
green	cross	→	?
blue	circle	→	?

Figure 2. Appearance-reality, conditional discrimination and configural learning presented in the same format.

Now I would like to draw attention to the correspondence between this task and a superficially very different task from an entirely different tradition, the conditional discrimination task, shown in Figure 2. In a typical conditional discrimination the correct response is dependent on background: For example, if the stimuli were triangle and square, on black or white backgrounds, configural discriminations might have the form:

background	cue		response
black	triangle	→	R+
black	square	→	R-
white	triangle	→	R-
white	square	→	R+

The background cue is analogous to the setting condition or conditional variable in appearance-reality tasks, the cue is analogous to the object properties, and the response is analogous to person's percept of the object. The two tasks are shown in correspondence in Figure 2.

Viewed this way it is easy to see that appearance-reality is like an incomplete conditional discrimination task. It is like a conditional discrimination task in which only one level of the object attribute is used. This suggests that we could vary appearance-reality tasks by varying the object attribute as well as the setting conditions.

Another thing to note is that conditional discrimination has proved very difficult for young children. It is one of the tasks that are usually called configural learning problems (Rudy, 1991). Each response is equally associated with both stimulus elements, and with both backgrounds. Thus R+ is associated with both triangle and square, and also with black and white. Therefore associative interference is high, and effectively blocks learning based on elemental associations.

An interesting point is that conditional discrimination is isomorphic to the exclusive-OR (XOR) and cannot be learned by two-layer neural nets, a fact that has had important implications for the theory of that field.

Conditional discrimination can be learned by fusing or chunking each stimulus into a unique configuration, thus:

Configuration triangle/black → R+

Configuration square/black → R-

Configuration triangle/white → R-

Configuration square/white → R+

The task can be learned this way because each configuration is distinct, and associative interference is reduced. The problem however is that, for each configuration to become unique, the elements must lose their identity. Thus if "Triangle" and "black" become fused into a unique configuration *Black- Triangle*, which is distinct from (say) *black-square* or *white-triangle*, then the components of black-Triangle lose their identity. An element within a configuration is, of necessity, no longer recognisable as the element it was. The problem then is that the structure of the task cannot be represented. The observable effect would be that the person would not be able to transfer to a problem isomorph.

Figure 2 shows an isomorphic conditional discrimination task. Notice that, once the first item is known, responses to the remaining 3 items can be predicted, and this would be true for any order of presentation. This can be done only if the relations in the task are represented. It is not possible if configurations are learned, so isomorphic transfer enables us to distinguish between configural learning and structural learning or learning of relations.

Transitivity and class inclusion

Transitivity and class inclusion are superficially dissimilar, to each other and to the other concepts we have been considering, yet at the relational level they share some important properties. The essence of transitivity is to integrate the premises into an ordered triple. For example, given the premises:

Bill is happier than Peter

Peter is happier than Tom

Most of us would arrange Bill, Peter, Tom from left to right, or top to bottom. Integrating the premises produces a slight experience of effort, confirmed experimentally, caused by the need to consider both premises jointly. For our present purposes the important thing about transitivity is order three elements. It entails processing a ternary relation.

Class inclusion entails representing the relations between the superordinate set B and the complementary subsets, A and A'. Given the problem in which apples and pears are included in fruit, the important thing is to identify which set is the superordinate. Fruit is not inherently a superordinate. Had the problem been fruit and meat included in food, fruit would have been a subordinate. Fruit is a superordinate because it includes two subordinate sets, apples and non-apples. Similarly, apples is a subordinate because it is included in fruit, and there is at least one other class that is included in fruit. The superordinate is identified by its relations to the other two sets. This entails processing a ternary relation.

Transitivity and class inclusion have a common degree of structural complexity, because the core of both is to process ternary relations. This imposes a high processing load for both children and adults. Relational complexity is a major cause of the difficulty that children have with both these concepts, and with many others that entail processing ternary relations.

In transitivity we have found that three- and four-year old children perform at near ceiling level when they only have to process one binary relation at a time. However when they have to integrate two binary relations into a ternary relation, their performance drops to approximately chance level. Other aspects of the task were held constant in these comparisons. A study of this kind is being presented by Glenda Andrews in this symposium.

$$\frac{1}{2} = \frac{3}{6}$$

$$\frac{1}{2} < \frac{4}{6}$$

$$\frac{1}{3} < \frac{3}{6}$$

$$\frac{5}{7} > \frac{5}{8}$$

Figure 3. Comparisons between fractions, illustrating that both numerators and both denominators must be processed, entailing representation interaction between four dimensions.

Fractions and proportions are known to cause difficulty for young children for reasons that never appear to have been completely explained. The problem is that proportion is a quaternary relation, and consequently imposes a high processing load, that is intrinsic to the concept. The examples in Figure 3 show that comparison of fractions entails working with four dimensions, because a change in either numerator or either denominator can affect the comparison. Comparison of integers requires working with only two dimensions. The powerful effect of structural complexity has been largely overlooked.

In order not to be misunderstood, there is nothing to prevent younger children from understanding fraction concepts that entail simpler relations. For example knowing what half a pie means entails only binary relations, and developmental data on processing relations indicates that it should be understandable in principle by two year olds. Notice that the predicted age is actually earlier than appears to have been observed. And while I am debunking myths, I would like to mention that good relational complexity analysis frequently uncovers unrecognised potential. For example it predicts that even two year olds should be able to make balance scale

judgments based on either weight or distance, though not both. This is more optimistic than current norms, and has been confirmed empirically by Cherie Dalton.

Processing relations by infrahuman animals

There do not appear to be many studies that have provided reasonably definitive evidence of infrahuman animals processing relations, and most tasks can be explained by associative processes. However Premack (1983) required chimpanzees to choose a pair of objects which had the same relation as a sample pair. In one variant, the sample comprised two objects that were the same, such as two apples. The participant had to choose another pair of objects that were the same, such as two bananas, in preference to a pair that were different, such as an apple and a banana. The task is a form of analogical reasoning, as Premack (1983) points out. The animals are required to select on the basis of the relation between a pair of objects, rather than by attributes of the objects. Only chimpanzees, and only those that had been language trained, could perform this task. It seems reasonable to conclude that chimpanzees can process binary relations, albeit only after extensive experience with symbols. However there appears to be no evidence that analogies based on binary relations can be attributed to lower animals.

PAIR MATCH-T0-SAMPLE

SAMPLE: XX

CHOICE*: YY XY

Analogy: O-same(XX) maps to O-same(YY)

XY

CHOICE*: YY AB

Analogy: O-different(XY) maps to O-different(AB)

*Correct choice underlined

Figure 4. Relational match-to-sample task.

These examples illustrate the point that structural complexity is a factor in conceptual difficulty. This not only true for children, but also for adults. Furthermore the structural complexity factor can potentially account for species differences among the higher animals, and for impairment resulting from frontal lobe damage.

I have tried to illustrate that analysis of relations entailed in tasks can reveal underlying properties. The common property in appearance-reality, conditional discrimination, transitivity and class inclusion is that they entail relating three variables. They are effectively ternary relations, and many tasks that have this structure are known to cause problems for young children.

Differences between associations and relations

Now I would like to consider the properties of relational knowledge and to contrast with the more basic architecture of associative knowledge. I will use the term *associative* in a way that is consistent with its most common usage, although I recognise that associations could be redefined to include some of the properties of relational knowledge. This however would obscure some important distinctions.

In relations the type of link can vary and is identified by a symbol (e.g. the link between "whale" and "dolphin" is explicitly symbolized by LARGER-THAN). This makes a relation accessible to other cognitive processes, so that a relational instance can be an argument to another relation. The property of using labelled links is shared by propositional networks, but is not characteristic of associations, in which all links are of the same kind, and unlabelled (Humphrey, 1951).

Higher-order relations have relations, or relational instances, as arguments, whereas first order relations have objects as arguments. For example: BECAUSE(AVOIDS(dolphin,whale), LARGER–THAN(whale,dolphin)). BECAUSE() is a higher-order relation

Associations can be chained, and can converge or diverge, but the associative link *per se* cannot be an entity in another association, so there is no associative equivalent of higher-order relations, and associations are not recursive.

Relational systematicity is possible because of higher-order relations (for example a>b → b<a). Implies is a higher-order relation.

Associations do not have this property.

Omni-directional access means that, given all but one of the components of a relation, we can access the remaining component. For example, given MOTHER-OF(woman,child), and given MOTHER-OF(woman,-) we can access "child", whereas given MOTHER-OF(-,child) we can access "woman", and given -(woman,child) we can access "MOTHER-OF".

Associations may be backward but omnidirectional access is not necessarily implied.

Representation of roles or slots, independent of specific instances. Thus BIGGER-THAN(-,-) entails representation of slots for a larger and a smaller entity.

Associations do not entail content-independent representations.

Dimensionality: each role defines in a dimension (an N-ary relation is a subset of the cartesian product of N sets). Thus a unary relation is a set of points in unidimensional space, a binary relation is a set of points in 2 dimensional space, and so on. This provides a conceptual complexity metric.

Associations are always reducible to sets of ordered pairs.

These properties mean that relational schemas have considerable generality and content independence. Because a relational link is represented by a symbol, and each role is explicitly represented, relational representations are independent of specific contents. For example, the relation *larger-than* remains the same irrespective of what arguments are bound to the symbol. Relations are the core of analogical reasoning, which enables relational schemas to be transferred across domains. We have also been to show how reasoning processes in tasks such as the Tower of Hanoi, and categorical syllogisms, can be attributed to relational representations. Relational schemas permit some powerful cognitive operations, and Steven Phillips has shown that operations on relational data bases can be implemented in neural net models of relational knowledge.

Cognitive Process	Neural net specification	Representational Rank
elemental association		0
configural association	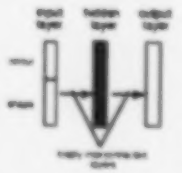	1
unary relation	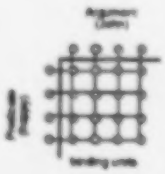	2
binary relation		3
ternary relation		4
quaternary relation		5
quinary relation	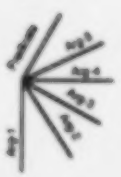	6

Figure 5. Representational ranks 0-6 with schematic neural net.

Relational knowledge and neural nets

The distinction between associative and relational knowledge corresponds to certain distinctions between neural net architectures. Two-layered nets can process elemental associations but not configural associations. Three-layered nets can process configural associations and can compute any computable function, but they do not typically embody the criteria for relational knowledge. The best way of representing relations in neural nets is currently a major topic of research, but one type of net that is suitable for this purpose is shown in Figure 5. The nets that represent relations have a representation for the relation symbol and a representation for each entity that is related. In our model the relation symbol and each argument is represented by vectors. A matrix of the vectors binds the representation. This is usually achieved by computing the tensor product of the vectors. Unlike three-layered nets, the components retain their identity. Thus the fundamental shift is that more than one dimension can be processed, while the dimensions retain their identity, whereas in three-layered nets the dimensions are fused. Access to this representation is omnidirectional and all the properties of relational knowledge are embodied.

Neural nets that represent relations incur a computational cost, and complexity analysis shows that this cost increases exponentially with number of entities related. This offers a natural explanation for the processing load observed in relational tasks, and for the limits to our processing capacity. I have also suggested that the phenomena that Piaget attributed to stages correspond in approximate fashion to levels of relational complexity as shown in Figure 5.

The balance scale is a good example of the contrast between relational and associative knowledge. It comprises four dimensions; weight and distance on the left and weight and distance on the right. Omni-directional access applies because input of any four components permits retrieval of the fifth; if we know both weights and both distances, we can infer whether the beam will balance. Alternatively, given weight and distance on the left, distance on the right, and the fact that the beam is balanced, it can decide what weight must be on the right. This is realistic because such tests are used in assessment (e.g. Surber & Gzesh, 1984), and because

we would be unwilling to attribute understanding to a child who could compute only one type of output. Thus a child who could tell us whether the beam would balance, but could say nothing about, for example, which weights and distances were required to produce a given state of balance or imbalance.

McClelland (1995) has shown that a three-layered net (with input, hidden and output layers of units) can be trained to indicate whether a beam will balance, given weight and distance on left and right as input. The model accounts for a number of important empirical observations which challenge earlier theories of children's balance scale performance. However this model does not fully represent the principle of the balance scale, and does not incorporate the omni-directional access property. Weight and distance on left and right must always be inputs in that model, and only one output, balance/left-down/right-down can be calculated. However a complete model of the balance scale can take *any* subset of N-1 variables as input and generate the Nth variable as output. Thus while McClelland's three-layered net efficiently computes a specific function of four variables, it does not represent the balance scale as a quaternary relation, and does not capture full understanding of the concept.

Empirical criteria for relational knowledge

We will briefly consider some of the basic empirical Criteria for Relational Knowledge.

Transfer between isomorphs should be greater than transfer between nonisomorphic tasks, with materials and procedure controlled.

Performance must be better than chance (frequently neglected).

To identify processing at a given level of relational complexity, performance must be compared with baseline for the next lower level of complexity (frequently neglected).

Domain knowledge must be adequate.

Participants must be fully trained in the procedure and must demonstrate perfect order performance on the task at the next lower level of complexity, with materials and procedure controlled.

Steps must be taken to ensure that the task cannot be performed by an alternative procedure or strategy that does not entail the relevant relation (one of the most common problems in developmental research).

Conclusion

The concept of relational schema captures many of the properties of higher cognitive processes. It can form the basis for analyses of the underlying structures of tasks, thereby revealing hidden commonalities. It is also the basis for a conceptual complexity metric that offers a new account of processing capacity limitations. This complexity metric accounts for previously unexplained sources of difficulty in many cognitive tasks, and also reveals some unrecognised potential in young children.

By defining the properties of relational schemas we also clarify the difference between higher cognitive processes and more basic associative processes. The differences have often been noted. For example, higher cognitive processes are slow, serial, effortful, flexible, content independent, explicit, cognitively accessible, and dynamic. Associative processes are fast, parallel, automatic, not under strategic control, implicit, not cognitively accessible, and dependent on input. Yet in much of our theorising we try to account for very different psychological phenomena with a single set of processes. The reason seems to be that we have not succeeded in defining the fundamental differences between levels of process. In order to do this we will have to give much more thought to the fundamental properties of relational schemas in cognition.

References

Frye, D., Zelazo, P. D., & Palfai, T. (1995). Theory of mind and rule-based reasoning. Cognitive Development, 10, 483-527.

Halford, G. S. (1993). Children's understanding: the development of mental models. Hillsdale, N. J.: Erlbaum.

Halford, G. S. (1996). Children's understanding of the mind: An instance of a general principle? [Review of The Child's Discovery of the Mind]. Contemporary Psychology, 41, 229-230.

Humphrey, G. (1951). Thinking: An introduction to its experimental psychology. London: Methuen.

McClelland, J. L. (1995). A connectionist perspective on knowledge and development. In T. Simon & G. S. Halford (Eds.), Developing Cognitive Competence: New Approaches to Cognitive Modelling. Hillsdale, NJ: Erlbaum.

Premack, D. (1983). The codes of man and beasts. The Behavioral and Brain Sciences, 6, 125-167.

Rudy, J. W. (1991). Elemental and configural associations, the hippocampus and development. Developmental Psychobiology, 24(4), 221-236.

Surber, C. F., & Gzesh, S. M. (1984). Reversible operations in the balance scale task. Journal of Experimental Child Psychology, 38, 254-274.